arXiv:1403.7084v2 [stat.ML] 2 Apr 2014

Constrained speaker linking

David A. van Leeuwen[1] *and Niko Brümmer*[2]

[1]Netherlands Forensic Institute, The Hague and Radboud University Nijmegen, The Netherlands
[2]AGNITIO Research, Somerset West, South Africa

Abstract

In this paper we study speaker linking (a.k.a. partitioning) given constraints of the distribution of speaker identities over speech recordings. Specifically, we show that the intractable partitioning problem becomes tractable when the constraints pre-partition the data in smaller cliques with non-overlapping speakers. The surprisingly common case where speakers in telephone conversations are known, but the assignment of channels to identities is unspecified, is treated in a Bayesian way. We show that for the Dutch CGN database, where this channel assignment task is at hand, a lightweight speaker recognition system can quite effectively solve the channel assignment problem, with 93 % of the cliques solved. We further show that the posterior distribution over channel assignment configurations is well calibrated.

Index terms: Speaker recognition, linking, partitioning, clustering, Bayesian.

1. Introduction

Speaker linking [1], also known as speaker partitioning [2], is the general problem of finding which utterances in a collection of recordings are spoken by the same speaker. It resembles speaker diarization [3], in that usually no prior training material of any of the speakers is given, but in speaker diarization there is the additional problem of segmentation, with the possibility of overlapping speech [4]. In [2] it is shown that the task covers many different problems in the area of speaker recognition, ranging from traditional 'NIST SRE-style' speaker detection to speaker counting and unsupervised adaptation. The authors approach the problem in a purely Bayesian way, meaning that the task is defined as computing a posterior distribution over all possible speaker partitions, given the data and a prior distribution over the partitions. Because of the combinatoric explosion of the number of partitions, the Bayesian approach can only be taken for small-sized problems. A different approach to essentially the same problem was taken in [1], where a solution was sought in terms of agglomerative clustering, i.e., making sequences of speaker linking decisions, concentrating on large scale problems. Speaker linking has found application in large scale speaker diarization [5–7], where the task of speaker diarization within a single recording is extended to finding the same speaker across multiple recordings, and to very long recordings for which the speaker diarization problem becomes computationally challenging [5].

In this paper, we are investigating how prior information in the speaker linking problem can be used. We will use a probabilistic approach following [2], and analyse prior information in terms of uncertainty. Then, we will apply this to a specific, but remarkably common situation in telephone speech data collections. This is the case where speaker identities in telephone conversations are known, but the assignment of identities to the two channels in the recorded conversation is unknown. We refer to this problem as the *channel assignment* task. In The Netherlands, we have seen this situation more than once. In the telephone interception recordings made in police investigations, both parties of a conversation are recorded in separate channels. The 'identity' of the speaker in such a case is limited to technical metadata such as the telephone number or the IMSI and IMEI numbers of the SIM and handset, respectively. This metadata is stored in the interception database, but for a reason unknown to us, it is impossible to tell which identity is recorded in which channel of the 2-channel audio file. From a forensic point of view, this does not pose a problem. A specific person is under investigation, and this means that calls to/from this person's telephone are intercepted. When in the content of the call there is incriminating evidence, it is that fragment that is going to be important. If the speaker identity of that fragment is questioned, forensic speaker comparison is going to play a role in the case. But even then, the channel assignment of the phone number is not important. This is different when we want to employ data mining approaches to all recordings in a police investigation. Then the link between identity and channel is very relevant indeed.

A second example is the data collection 'Dutch Spoken Corpus' (*Corpus Gesproken Nederlands*, CGN [8]). This is a large general speech database of contemporary Dutch spoken in the Netherlands and Belgium, with a wide variety of speech sources and speaking styles, and annotated at various levels of details. As such, it is widely used by researchers in linguistics, language and speech technology in these countries. In 2008, the data was used as training material in the project N-Best, an evaluation of Dutch speech recognition systems [9]. At the preparation stage it turned out that, despite the multitude of annotations available with CGN, the mapping of orthographic annotation (and speaker identities) to channels within a telephone conversation was unknown. A possible reason for this is that the telephone conversations recorded as two channels only contribute to a small portion of the entire CGN, and that despite the very broad application scenario included in the design of CGN, requirements for automatic speaker recognition may not have been fully worked out. At the time, additional manual speaker attribution was made and distributed to the participants of N-Best. Participants reported different strategies for dealing with the originally missing speaker-to-channel annotation [10, 11].

Other situations where speaker identities are known, but the exact mapping is unknown are large scale diarization problems. E.g., in [6], a sequence of meetings is processed with diarization, and later the speaker linking between meetings is carried out. Sometimes the speaker information is partially known, as in the case of the speaker diarization of broadcast shows where TV-guide metadata can reveal the names of some of the speakers in the shows [12].

2. Reducing uncertainty

We are going to describe the speaker linking problem in terms of uncertainty, where both prior information and speaker recognition can contribute to lowering the uncertainty in the speaker identities in the database. We will do this in terms of the total entropy in the database, parameterized by size and other prior information. The number of recordings will be denoted by $2M$, anticipating that we are going to investigate a collection of M conversations, each contributing two separate recordings, shortly. In the following, we will consecutively add constraints, or prior information, to the analysis.

2.1. Speaker-homogeneous recordings

The overall restriction in this paper is that each speech recording contains speech from only a single speaker. This excludes the tasks of speaker diarization. The speaker entropy is $H_{\rm I} = \log B_{2M}$, where B_n is the nth 'Bell number' [2]. Here we have applied no further priors to the number of speakers and their distribution, not even that—for a telephone call—a speaker can't talk to herself.

2.2. Number of speakers is known

When additionally the number of speakers N in the database is known, the entropy is reduced to $H_{\rm II} = \log \left\{ {2M \atop N} \right\}$, where $\left\{ {n \atop k} \right\}$ denotes the Stirling number of the second kind. An upper bound approximation is $2M \log N - \log N!$, i.e., for each recording the identity can be any of N speakers, and we must compensate for the arbitrary speaker labelling.

2.3. Telephone conversations

When we know that the conversations are all telephone conversations, we can exclude situations where one speaker occurs in both sides of the conversation. This reduces the entropy just a little further by approximately $M \log \frac{N}{N-1}$.

2.4. Speakers in conversations are known

We now make a big step to the channel assignment task described in the introduction: the speaker identities of the parties participating in a conversation are known, but it is unknown in which channel which speakers is. This prior results in the reduction of entropy to $H_{\rm IV} = M \log 2$. If the entropy is expressed in bits, we have $H_{\rm IV} = M$.

2.5. Application of speaker recognition

Speaker linking can further reduce the entropy. For instance, if a single speaker occurs in all M conversations, and there are $M + 1$ speakers in total, it is easy to see that, if speaker recognition works flawlessly, the entropy can be reduced to 0. However, there are many different partitionings of the speakers over the calls possible, and not all will have te same potential w.r.t. entropy reduction even with a perfect speaker recognition system. For instance, if the same two speakers occur in all M conversations, there is still $\log 2$ entropy left. And in the extreme situation that there are $2M$ speakers in the database, the entropy remains at $M \log 2$ as before.

The entropies described in the previous paragraph are *potentially attainable* entropies in the case of perfect speaker linking. However, speaker recognition is based on statistical models of speech and speakers, and can make errors. We therefore need a different measure to compute the effect of speaker recognition, and we will use the cross entropy for that.

Table 1: Some example values for the entropies of the speaker paritioning problem. We have used the parameters of the test data from CGN (348 conversations involving 356 speakers) as an example of the reduction in entropy H, which is expressed in bits here. F is the average confusion, see Section 4.1.

Section	Expression H	CGN H	CGN F
2.1	$H_{\rm I} = \log B_{2M}$	4163.2	3991
2.2	$H_{\rm II} = \log \left\{ {2M \atop N} \right\}$	3292.7	704
2.3	$H_{\rm III} = H_{\rm II} - M \log \frac{N}{N-1}$	3291.3	702
2.4	$H_{\rm IV} = M \log 2$	348	1.0

3. Speaker recognition

We will use de detection capability of a speaker recognition system as the information used in speaker linking. We assume that when two utterances (different conversation sides) x and y are available, the recognizer can provide a log-likelihood-ratio for comparing the two

$$\lambda(x, y) = \log \frac{P(x, y \mid \mathcal{H}_1)}{P(x, y \mid \mathcal{H}_2)} \tag{1}$$

where $\mathcal{H}_1$ and $\mathcal{H}_2$ are the hypotheses that x and y are spoken by the same or different speakers, respectively. We assume that the system is *well calibrated*, i.e., that λ can be used effectively to compute a minimum risk Bayes decision.

We will utilize the speaker detector for analyzing various sub-partitionings of the telephone channel assignment problem. We will have to introduce some notation at this point. A *conversation labelling* is the way speaker pairs are distributed over telephone conversations, i.e., without explicitly encoding which speaker is in which channel. If such a conversation labelling encompasses M conversations, there are 2^M possible configurations $\mathcal{C}$ of the speaker pairs over the channels in the conversations. These can be numbered in binary notation, using ℓ and r instead of traditional 0's and 1's. The channel assignment problem of Section 2.4 is to find the correct configuration given the conversation labelling and the speech data.

For a given configuration $\mathcal{C}$ of speakers over recordings, the likelihood is $L_{\mathcal{C}} = P(X \mid \mathcal{C})$, where X denotes the relevant speech. Given the prior probability for this configuration $\pi_{\mathcal{C}}$ and the likelihoods, the posterior can be computed using

$$P(\mathcal{C} \mid X) = \frac{\pi_{\mathcal{C}} L_{\mathcal{C}}}{\sum_{\mathcal{C}'} \pi_{\mathcal{C}'} L_{\mathcal{C}'}}, \tag{2}$$

where the summation is over all 2^M possible configurations $\mathcal{C}'$. For the stated task, it is not unreasonable to set the prior uniformly at $\pi_i = 2^{-M}$.

In Table 1 we have summarized the way the reduction of entropy evolves by specifying more constraints. As a numerical example, we have used the test data in the experiment explained in Section 4.

3.1. Single speaker chains

Let us first consider the simplest case, a set of two conversations with a single common speaker. We denote this 'target speaker' by a, occurring twice, with conversation partners b and c, respectively. There are four possible configurations over the two channels ℓ, r for the conversations 1 and 2, namely (ab, ac), (ab, ca), (ba, ac) and (ba, ca), cf. Fig. 1. Using $L_{\rm norm}$ to denote the likelihood that all recordings have different identities,

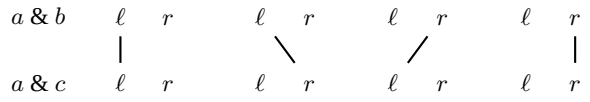

Figure 1: The four linking configurations for a single speaker chain in 2 conversations. The link indicates the channel of a.

the log-likelihood for the first in the four configurations is

$$\log L_1 = \log L_{\text{norm}} + \lambda(1\ell, 2\ell), \tag{3}$$

with similar likelihoods $L_2, \ldots, L_4$ for the other configurations. Note that with (2), the factor L_{norm} cancels in the posteriors P_i. A maximum posterior linking decision would therefore correspond to a clustering step based on maximum likelihood.

Next, we will consider a situation with M conversations, each with the same 'target speaker' a, and all different conversation partners. Using L_{norm}^{2M} to denote the likelihood that all $2M$ recordings are from different speakers, a first approximation to the likelihood $\tilde{L}_1$ for the first configuration $(ab, ac, \ldots, aM)$ could be computed using

$$\log \tilde{L}_1 = \log L_{\text{norm}}^{2M} + \sum_{i=1}^{M-1} \lambda(i\ell, (i+1)\ell), \tag{4}$$

i.e., making the "link" between assumed speakers a in channel ℓ from 1 to 2, from 2 to 3, etc. But this is just one possible linking, and there are in fact $\binom{M}{2}$ possible log-likelihood-ratios, from which a chain of $M-1$ need to be chosen. A better approximation therefore is to include them all, and scale these to $M-1$ contributions:

$$\log L_1 = \log L_{\text{norm}}^{2M} + \frac{2}{M} \sum_{i<j} \lambda(i\ell, j\ell). \tag{5}$$

3.2. Re-occurring conversation partners

Thus far we have considered that case that all conversation partners of a particular target a are different. However, it is likely that in some conversations the same partner b occurs. This leads to two different cases:

unresolvable Target a and b are always paired in the database. In this case, a relative linking is possible, but the absolute attribution of the speakers to the links cannot be resolved.

resolvable There is more than one conversation partner for a (or b). In principle, the speaker linking could be resolved by speaker recognition.

If we further concentrate on the subset of the database involving speaker a, then multiple occurrences of b can add to the log-likelihood. The additional contribution of speaker b to the log-likelihood of the 1st configuration is, similar to (5)

$$\frac{2}{M_b} \sum_{i<j} \lambda(ir, jr) \tag{6}$$

where the sum is over all M_b conversations involving b as a conversation partner of a—for the first configuration b is in channel r. This average log-likelihood ratio should be added to L_1 of (5), and similar terms to the log-likelihoods of the other configurations.

3.3. Independent cliques

If there are two sets of speakers $\mathcal{S}_1$ and $\mathcal{S}_2$ in the database that do not share a single conversation, we call these different *cliques*. When $\mathcal{C}_1$ is a configuration within all conversations X_1 involving $\mathcal{S}_1$, and $\mathcal{C}_2$ correspondingly for $\mathcal{S}_2$, then the posterior

$$P(\mathcal{C}_1, \mathcal{C}_2 \mid X_1, X_2) = P(\mathcal{C}_1 \mid X_1) P(\mathcal{C}_2 \mid X_2), \tag{7}$$

i.e., the configurations can be treated independently, reducing the total number of configurations that need to be computed. If the design of the database does not allow for the analysis of separate cliques, the problem of computing the posterior distribution over all configurations of the database will still be intractable for large M.

4. Experiments

In this section we apply the constrained speaker linking approach sketched above to 'NL/component c' of the Dutch CGN database [8], which consists of telephone conversations between acquaintances in The Netherlands, recorded using a telephony platform. As mentioned before this database actually lacks the channel-to-speaker assignment, but within the context of the N-Best project [9] this information had been manually added. We treat this part of CGN as a prototypical example of the channel assignment task, and use the manually added reference for evaluation purposes. It consists of 352 conversations (704 sides) involving 357 different speakers.

The speaker recognition system we use is a lightweight 'UBM/GMM dot-scoring' system [13] implemented entirely in the new high performance language for numerical computing Julia.[1] We use a standard acoustical front end based on 20 MFCC's plus first and second derivatives, energy-based Speech Activity Detection, and 4 second feature warping [14]. The 1024 component UBM was trained gender-independently on 'NL/component d' (telephone conversations recorded using a mini-disc, 600 conversations, 176 speakers, 21 hours). The speaker comparison score is a linear approximation to the MAP-adapted [15, 16] GMM / UBM log likelihood ratio score. For probabilistically interpretable likelihood ratios (1) we self-calibrated the collection of test scores using CMLG [17] before applying (2).[2] The equal error rate of the system evaluated on a full score matrix of all recordings of the test data is $E_= = 7.0\,\%$.

NL/component c of CGN consist of many independent cliques, presumably as a result of way the speakers were recruited. There are 125 cliques of 2–5 speakers, each having 1–6 conversations amongst each other. During the course of this study, the recognition system hinted on potential labelling errors. We manually listened to the conversations of cliques to check the labelling. We started with the largest cliques of 6 conversations, selecting only those with the biggest log error, and worked our way down in clique size. In most cases labeling errors were very clear, observable from gender, mentioned names and relations, as many cliques seemed to have been formed within families. We corrected the labeling of suspicious cliques once, without further feedback from the speaker recognition and linking system. As a result of various dubious speaker labels, we ended up using 348 conversations involving 356 speakers.

[1] `http://julialang.org/`

[2] This procedure is comparable to determining $C_{\text{det}}^{\text{min}}$ in speaker detection evaluation.

Table 2: Performance of the linking experiment on CGN, Dutch, component c, grouped by complexity. $|\mathcal{C}|$ is the number of configurations for a clique, $N_{\mathcal{C}}$ is the number of cliques of this complexity. Entropies are measured in bits.

$\lvert\mathcal{C}\rvert = 2^{M_{\mathcal{S}}}$	$N_{\mathcal{C}}$	$\overline{H}_{\text{cross}}$	F	E (%)
2	8	1.0	1.0	NA
4	61	0.205	0.153	3.3
8	13	0.201	0.150	7.7
16	26	0.032	0.023	7.7
32	9	0.041	0.029	22
64	5	0.019	0.013	20
all resolvable	110	0.078	0.056	7.0

4.1. Evaluation metrics

A natural metric that extends the earlier mentioned entropy considerations is the cross entropy. The cross entropy between the evaluator's and recognizer's posteriors for a clique $\mathcal{S}$ with a true partitioning $\mathcal{P}$ is

$$H_{\text{cross}} = -\sum_{\mathcal{C}\in\mathcal{S}} P(\mathcal{C} \mid \mathcal{P}) \log P(\mathcal{C} \mid X) \tag{8}$$

$$= -\log P(\mathcal{P} \mid X), \tag{9}$$

i.e., H_{cross} can be seen as a logarithmic scoring rule, which is known to be strictly proper [18]. Because H_{cross} covers $M_{\mathcal{S}}$ conversations, the average cross entropy per conversation is

$$\overline{H}_{\text{cross}} = \frac{H_{\text{cross}}}{M_{\mathcal{S}}}. \tag{10}$$

An intuitive meaning of this average entropy is through the perplexity, $\exp \overline{H}_{\text{cross}}$, as the average number of configurations to choose from for a set of conversations. Along the lines of the Albayzin Language Recognition Evaluation metric [19], we will rather use the *confusion*

$$F = e^{\overline{H}_{\text{cross}}} - 1. \tag{11}$$

It represents the average number of *wrong* alternatives. F is 0 for a system operating perfectly. In Table 1 we have tabulated the reduction of F based on the prior entropy. In the last row, which represents our channel assignment task, there is just one configuration alternative to the correct one. The goal is to further reduce this to 0 by applying speaker linking.

Finally, we define the clique error rate E as the average number of cliques for which the maximum posterior configuration is not the true speaker configuration.

4.2. Linking

In Table 2 the results of the speaker attribution experiment for CGN are shown. We have conditioned the performance on the complexity $|\mathcal{C}| = 2^{M_{\mathcal{S}}}$, the number of configurations in the speaker attribution task, thereby averaging over cliques with the same number of configurations. For a clique of one conversation, $|\mathcal{C}| = 2$, speaker linking cannot resolve the uncertainty, $P_{1,2} = 0.5$, and the entropy is 1 bit. There are 8 such speaker pairs that both occur only once in the database, and a further four cases with $|\mathcal{C}| = 4$ where the same speakers are paired twice. We will further discard these unresolvable cases in the analysis.

The average cross entropy drops below 1 bit for all complexities above 2, which is an increasing reduction w.r.t. the original entropy $\bar{H} = 1$. The confusion drops steadily with increasing complexity, and finally, partitioning error rate is more or less stable at about 7 % on average. We have computed $N_{\mathcal{C}}$-weighted averages over the resolvable cliques in the last row of the table.

Table 3: The effect of calibration. $\overline{H}_{\text{cross}}$ is the average cross entropy before re-calibration, cf Table 2. The data are taken only over resolvable cliques.

$M_{\mathcal{S}}$	$\overline{H}_{\text{cross}}$	$\overline{H}^{\min}_{\text{cross}}$	$a_{\min}/a$
2	0.082	0.071	1.48
3	0.201	0.177	0.65
4	0.032	0.029	1.42
5	0.041	0.040	0.77
6	0.019	0.018	0.79
all	0.078	0.076	1.17

4.3. Calibration

The scores λ have been 'self-calibrated' using a linear transform $\lambda = as + b$, where s is the score and a and b are parameters found to minimize the 2-class cross entropy (e.g., C_{llr}) in a classical detection set-up, quite similar to logistic regression. For the posterior (2) the offset b cancels, so effectively a is the only parameter of importance in this task.

We want to investigate if the approximation by averaging of log likelihoods as in (5) and the addition of independent evidence as in (6) is not making the system over-confident. Therefore we re-calibrate the scaling factor a for different clique complexities, and investigate the change in a and the improvement in average cross entropy per conversation.

In Table 3 the improved $\overline{H}^{\min}_{\text{cross}}$ for the different complexities is shown, with the additional factor $a_{\min}/a$ applied to the recognizer's log likelihoods. These values of $a_{\min}/a$ vary a bit, but given the relatively small numbers of cliques that are involved in the optimization, this may be expected. The optimization over all resolvable cliques results in a value $a_{\min}/a = 1.17$, which is very close to unity. This means that the original calibration based on detection trials only was good, and that the averaging operation (5) and addition of evidence from other links in the clique (6) do not make the posteriors overconfident.

5. Conclusions

We have seen in the channel assignment task that with higher complexity, i.e., higher prior entropy, the average cross entropy per conversation quickly drops towards zero. On the one hand, with increasing complexity the task gets harder because there are more configurations to discriminate between. On the other hand, with higher complexity the increased clique size provides more opportunity for the recognition system to determine the correct configuration with confidence. These two effects more-or-less cancel in the fraction of correctly found configurations. With a very lightweight speaker recognizer, trained with a small amount of domain specific speech material, the channel assignment task can be carried out with a high degree of success. This is indeed a complete Bayesian solution to the speaker partitioning problem as proposed in [2], which is deemed intractable in general. In cases like the CGN database, where conversation clique sizes $M_{\mathcal{S}}$ are relatively small, the computational load still is negligible (computing the posteriors for all cliques, given all relevant log-likelihood-ratios, takes about 1 second), but of course the load grows exponentially with $M_{\mathcal{S}}$. However, it is likely that also in forensic investigations the cliques are small so that we expect that our approach can be used there as well.

6. References

[1] D. A. van Leeuwen, "Speaker linking in large data sets," in *Proc. Odyssey*. IEEE, June 2010, pp. 202–208.

[2] N. Brümmer and E. de Villiers, "The speaker partitioning problem," in *Proc. Odyssey 2010: The Speaker and Language Recognition Workshop*, Brno, June 2010.

[3] S. Tranter and D. A. Reynolds, "An overview of automatic speaker diarization systems," *IEEE Trans. on Audio, Speech and Language Processing*, vol. 14, no. 5, pp. 1557–1565, 2006.

[4] M. Huijbregts, D. van Leeuwen, and F. de Jong, "Speech overlap detection in a two-pass speaker diarization system," in *Proc. Interspeech*. Brighton: ISCA, September 2009, pp. 1063–1066.

[5] M. Huijbregts and D. A. van Leeuwen, "Large scale speaker diarization for long recordings and small collections," *Audio, Speech, and Language Processing, IEEE Transactions on*, vol. 20, no. 2, pp. 404–413, 2012.

[6] M. Ferras and H. Bourlard, "Speaker diarization and linking of large corpora," in *Proc. Spoken Language Technology Workshop*, Miami, FL, 2012, pp. 280–285.

[7] W. M. Campbell and E. Singer, "Query-by-example using speaker content graphs," in *Proc. Interspeech*, Portland, 2012, pp. 1095–1098.

[8] N. H. J. Oostdijk and D. Broeder, "The Spoken Dutch Corpus and its exploitation environment," in *Proceedings of the 4th International Workshop on Linguistically Interpreted Corpora (LINC-03).*, Budapest, Hungary, 2003.

[9] D. A. van Leeuwen, J. Kessens, E. Sanders, and H. van den Heuvel, "Results of the N-Best 2008 Dutch speech recognition evaluation," in *Proc. Interspeech*. Brighton: ISCA, September 2009, pp. 2571–2574.

[10] J. Despres, P. Fousek, J.-L. Gauvain, S. Gay, Y. Josse, L. Lamel, and A. Messaoudi, "Modeling Northern and Southern varieties of Dutch for STT," in *Proc. Interspeech*. Brighton: ISCA, 2009, pp. 96–99.

[11] M. Huijbregts, R. Ordelman, L. Werff, and F. Jong, "SHoUT, the University of Twente submission to the N-Best 2008 speech recognition evaluation for Dutch," in *Proc. Interspeech*. ISCA, 2009, pp. 2575–2578.

[12] M. Huijbregts and D. A. van Leeuwen, "Diarization-based speaker retrieval for broadcast television archives," in *Proc. Interspeech*. Firenze: ISCA, August 2011.

[13] A. Strasheim and N. Brümmer, "SUNSDV system description: NIST SRE 2008," NIST SRE workshop proceedings, May 2008.

[14] J. Pelecanos and S. Sridharan, "Feature warping for robust speaker verification," in *Proc. Speaker Odyssey*. Crete, Greece, 2001, pp. 213–218.

[15] J.-L. Gauvain and C.-H. Lee, "Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains," *IEEE Trans. Speech Audio Processing*, vol. 2, pp. 291–298, 1994.

[16] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted gaussian mixture models," *Digital Signal Processing*, vol. 10, pp. 19–41, 2000.

[17] D. A. van Leeuwen and N. Brümmer, "The distribution of calibrated likelihood-ratios in speaker recognition," in *Proc. Interspeech*. ISCA, 2013, pp. 1619–1623.

[18] M. DeGroot and S. Fienberg, "The comparison and evaluation of forecasters," *The Statistician*, pp. 12–22, 1983.

[19] L. J. Rodríguez-Fuentes, N. Brümmer, M. Penagarikano, A. Varona, G. Bordel, and M. Diez, "The Albayzin 2012 language recognition evaluation," in *Proc. Interspeech*, 2013, pp. 1497–1501.